

WHAT IS CLAIMED IS:

1. A method for detecting and repairing a cloud partition in a network, wherein a cloud includes a plurality of nodes, each node having a peer identifier (ID) and a cache of peer IDs for one or more known nodes, the method comprising:

requesting identification information for an adjacent node from a server node;
receiving identification information for a closest-match node that is a closest found match for the adjacent node; and

if a peer ID for the closest-match node is not in the cache, determining that the cloud is partitioned, and sending a repair message to a known node indicating the identification information for the closest-match node.

2. The method of claim 1 wherein the peer ID is numeric, and wherein a first node having a first peer ID is adjacent to a second node having a second peer ID if there are no nodes using a peer ID that is between the first peer ID and the second peer ID.

3. The method of claim 1 wherein the cache is a multilevel cache, and wherein a host node records an entry for a known node in a level of the multilevel cache in accordance with a distance between a peer ID of the known node and a peer ID of the host node, and wherein entries for known nodes closest to the host node are recorded in a lowest level of the cache.

4. The method of claim 3 further comprising flooding the identification information of the closest-match node to nodes having entries in the lowest cache level, and flooding identification information of the host node to the closest-match node.

5. The method of claim 4 further comprising receiving a flood message from a node in a different partition of the cloud.

6. The method of claim 4 further comprising determining that the closest-match node is in a larger partition when the closest-match node is closer than every node in the cache, and

sending each node in a top level of the cache a repair message including a peer ID and an address for a node in the larger partition, and a repair level indicating repair is required for a cache level immediately below the top level.

7. The method of claim 1 further comprising determining that the closest-match node is in a smaller partition when the closest-match node is not closer than every node in the cache, and sending a repair message to the closest-match node including identification information for a known node.

8. A computer-readable medium having computer-executable instructions for performing the steps of claim 1.

9. A method for detecting and repairing a cloud partition in a network, wherein a cloud includes a plurality of nodes, each node having a peer identifier (ID) and a cache of peer IDs for one or more known nodes, the method comprising:

receiving a repair message at a host node including identification information for a first node;

requesting from the first node identification information for a second node that is adjacent to the host node;

receiving identification information for a closest-match node that is a closest found match for the second node; and

if a peer ID for the closest-match node is not in the cache, determining that the cloud is partitioned, and sending a repair message to a known node indicating the identification information for the closest-match node.

10. The method of claim 9 wherein the peer ID is numeric, and wherein the second node having a second peer ID is adjacent to the host node having a host peer ID if there are no other nodes using a peer ID that is between the second peer ID and the host peer ID.

11. The method of claim 9 wherein the cache is a multilevel cache, and wherein a host node records an entry for a known node in a level of the multilevel cache in accordance with a distance between a peer ID of the known node and a peer ID of the host node, and wherein entries for known nodes closest to the host node are recorded in a lowest level of the cache.

12. The method of claim 11 further comprising flooding the identification information of the closest-match node to nodes having entries in the lowest cache level, and flooding identification information of the host node to the closest-match node.

13. The method of claim 12 further comprising receiving a flood message from a node in a different partition of the cloud.

14. The method of claim 12, wherein the repair message further includes a repair level L , the method further comprising determining that the closest-match node is in a larger partition when the closest-match node is closer than every node in the cache, and sending each node in a level of the cache equal to L a repair message including a peer ID and an address for a node in the larger partition, and a repair level equal to $L+1$.

15. The method of claim 9 further comprising determining that the closest-match node is in a smaller partition when the closest-match node is not closer than every node in the cache, and sending a repair message to the closest-match node including identification information for a known node.

16. A computer-readable medium having computer-executable instructions for performing the steps of claim 9.

17. A method for detecting a cloud partition in a network, the method comprising:
estimating a size of the cloud; and
testing for a cloud partition at a time interval inversely proportional to the estimate of the size of the cloud.

18. The method of claim 17 wherein estimating a size comprises:
establishing a multilevel cache of node entries wherein node entries are recorded in levels of the multilevel cache based on distance from a cache establisher, and wherein the quantity of levels in the cache is L and a ratio of a number space covered by each level is K ; and
computing an estimate of the size of the cloud from L and K .
19. The method of claim 18 wherein a distance between the cache establisher and a second node is measured as a distance between a peer ID of the cache establisher and a peer ID of the second node.
20. A computer-readable medium having computer-executable instructions for performing the steps of claim 17.
21. A computer readable medium containing a repair message data structure for initiating repair of a cloud split in a network, the message data structure comprising:
a first data field representing a peer identifier of a node to locate;
a second data field representing an address of a node to be queried; and
a third data field representing a cache level to be used for a repair test.